3SC

The Decision Gap in Pharma Supply Chains

Why Visibility Has Improved Faster Than the Ability to Act

SEPTEMBER 2026



CONTENTS

01	Executive summary
02	Complexity has outgrown functional visibility
03	The real constraint is decision latency
04	The control model: one decision loop
05	Digital Control Tower
06	Risk Intelligence
07	Agentic AI
08	Decision Governance
09	Provenance and Trust
10	The adoption reality check
11	How the ecosystem works together
12	Transport Execution
13	Where the value sits
14	Closing Perspective

EXECUTIVE SUMMARY

Most pharma companies can tell you where a shipment is. Far fewer can tell you on the day it matters, whether the batch on board will still be sellable when it lands, which customers it is still eligible to serve, who is authorised to redirect it, and what the best alternative is worth.

The second question is where the money goes.

According to Gartner's life science and pharmaceutical supply chain research, fewer than half of life science supply chains use technology to model how different scenarios would affect the outcome of a decision. Most are still working on the trade-offs by hand, in meetings, after the exception has already been escalated.

Gartner attributes this to siloed organisations, poor access to data, and legacy technology. None of those are detection failures. The signal usually arrives but in a system that cannot say what to do about it, to a team that cannot approve the fix on its own.

Meanwhile the exposure keeps growing. Pharmaceutical logistics is forecast to expand from roughly US$108 billion in 2025 to US$208 billion by 2033 as per Grand View Research, with cold chain expected to be the fastest-growing segment. This means more pharmaceutical products will be moving across increasingly complex networks, often under tighter temperature and shelf-life constraints and through partners and infrastructure that enterprises do not directly control.

Yet many pharma operating models still treat risk as a functional exception. Inventory teams manage stock. Logistics teams manage shipments. Quality teams manage release and excursions. Planners manage demand and supply. Commercial teams manage customer commitments. Finance sees the economic impact once the decision has already been made.

The disruption is almost always visible to someone on that list. The decision requires all of them. That is the gap, and it is measured in days the company did not have.

This paper argues that the next maturity step is not another visibility dashboard. It is a move from visibility-led control towers to execution-led control models: one decision loop that connects product, batch, inventory, shipment, supplier, quality, customer, and external disruption signals, so that risk becomes a governed decision while options still exist.

The practical opportunity:

- Fewer expiry losses, because stock is managed by serviceability rather than availability
- Stronger cold-chain control, because lane risk is detected before excursion
- Faster rerouting, because options are compared across cost, time, quality, and service
- Lower logistics leakage, because premium freight, demurrage, and detention are governed earlier
- Higher planner productivity, because decision context is assembled before escalation
- Better executive control, because operational risk is translated into value exposure

The scale is visible in adjacent industries. In McKinsey's January 2026 analysis of manufacturing footprints, an aerospace supplier that rebuilt its capacity model around material availability and supplier performance cut expedited-service costs by 30% to 50%.

It did not negotiate cheaper expedites. It stopped needing so many of them.

Pharma runs on the same mechanism, with shelf life, cold chain, and quality release added on top, and with far less room for error.

COMPLEXITY HAS OUTGROWN FUNCTIONAL VISIBILITY

The traditional pharma supply chain was built around reliability, compliance, and product protection. Those priorities remain. What has changed is the operating complexity around them.

Growth is shifting toward biologics, specialty therapies, outsourced networks, temperature-sensitive distribution, and regionalised supply. McKinsey's January 2026 analysis places life sciences in the medium-disruption band, but with a warning attached. Life sciences and chemicals were historically more insulated from supply chain pressure than automotive or aerospace. The recent wave of disruption, spanning more countries and deeper supplier tiers than previous cycles, is exposing vulnerabilities in those sectors for the first time.

Cost pressure compounds it. Gartner notes that the healthcare supply chain cost to serve averages 37.3% of the total cost of patient care. Supply chain is not a back-office line item in this industry. It is roughly a third of what it costs to treat a patient.

That matters because the most material pharma risks are no longer contained inside one function.

A delayed inbound shipment is not just a logistics event. It may affect quality release, cold-chain exposure, market allocation, customer orders, expiry risk, and working capital.

A short-dated batch is not just an inventory event. It may reflect demand volatility, poor node positioning, slow release, weak allocation rules, or transport lead-time risk.

A port delay is not just dwell time. It may create demurrage, reefer exposure, emergency transfer, premium freight, service failure, and write-off risk.

A reroute is not just a recovery action. It is a decision across cost, service, lane qualification, temperature integrity, documentation, customer priority, and batch viability.

Exhibit 1: Complexity is expanding on three fronts

Complexity front	What is changing	Why it matters operationally
Product complexity	More specialty, biologic, and temperature-sensitive products	Stricter shelf-life, temperature, release, and handling controls
Network complexity	More outsourcing, cross-border movement, CMOs, 3PLs, and regional nodes	More risk sits outside direct enterprise control
Decision complexity	Faster disruptions, tighter service windows, higher cost of late action	Teams need earlier trade-off decisions, not later reporting

The operating challenge is no longer whether the enterprise has data. It is whether that data is organized around the decisions that protect service, product integrity, and margin.

THE REAL CONSTRAINT IS DECISION LATENCY

Most pharma companies have enough data to detect exceptions. They do not have enough connected context to decide early.

Shipment data sits in a freight portal. Inventory sits in ERP. Batch status sits in a quality system. Demand sits in a planning tool. Port congestion arrives from external feeds. Temperature data sits in a logger, an IoT platform, or a 3PL system.

The business question cuts across all of them:

What is exposed, when will it hit, which options still exist, who must approve the action, and how much value can be protected?

That question is rarely answered by a single system, and even more rarely with enough visibility across the wider network to support an early decision. McKinsey's supply chain risk pulse, published in December 2025, found that while nearly all companies surveyed can see risk at their tier-one suppliers, only 42% can see past them. For a pharma network built on CMOs, contract packagers, and specialist logistics providers, that is the boundary at which most external events stop being traceable to the flows they will actually disrupt.

The cost shows up in familiar ways. Planners discover near-expiry risk after the transfer window has closed. Logistics teams approve premium freight before checking alternate node inventory. Quality teams receive incomplete shipment context after a cold-chain exception. Supply chain heads learn that service is exposed only after local teams escalate. Finance sees freight leakage, write-offs, penalties, demurrage, and detention after they are locked in.

Decision latency is not a technology problem alone. It is an operating-model problem. It occurs when data, ownership, workflows, and execution rights are not aligned around risk.

Exhibit 2: Disruptions become expensive when teams see fragments instead of the cascade

Event	Functional view	Enterprise risk view
Port dwell increases	Shipment delay	Cold-chain exposure, demurrage, stockout, order risk
Batch approaches expiry	Aging inventory	Recoverability, market eligibility, customer acceptance, write-off risk
Carrier misses milestone	Transport exception	OTIF exposure, temperature risk, alternate routing need
Demand shifts by market	Planning variance	Wrong-node inventory, serviceable-stock gap, emergency transfer
Supplier or CMO delay	Procurement update	Production impact, order failure, allocation conflict

The more fragmented the view, the later the decision. The later the decision, the narrower and more expensive the options become.

THE CONTROL MODEL: ONE DECISION LOOP

A mature pharma control model does not begin with a dashboard. It begins with a decision loop. The purpose is simple: reduce the time between a material signal and a governed response.

That loop has six moves.

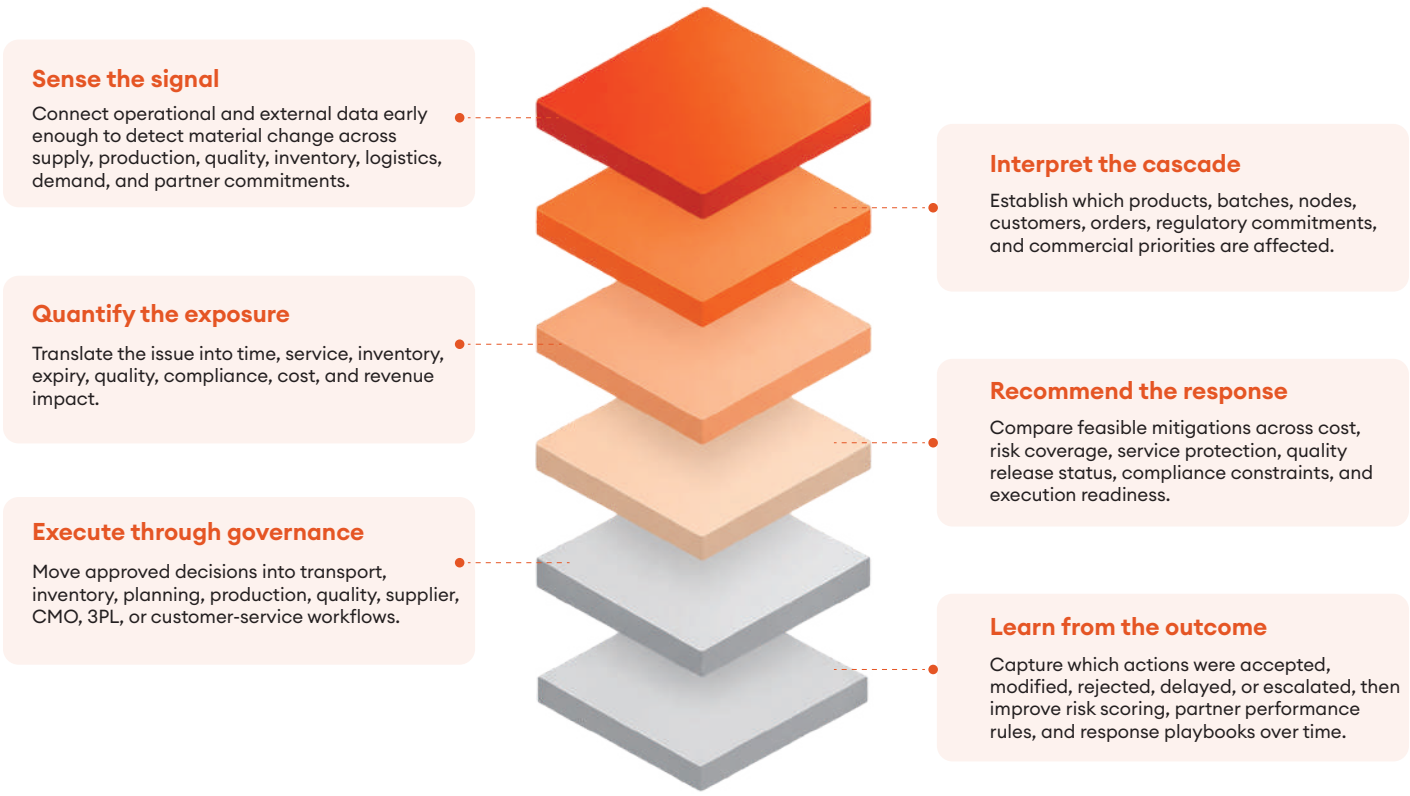


Figure: The six-move decision loop

The ecosystem works when the loop is not trapped inside one function. Planning sees execution risk. Logistics sees service and inventory exposure. Quality sees the downstream consequence of release delays. Commercial teams see what customer commitments are truly at risk. Leadership sees the few decisions that materially change the outcome.

Four capability layers make that loop work.

DIGITAL CONTROL TOWER

The control tower is the shared decision environment. Not another layer of dashboards.

In pharma, visibility is only useful when it is connected to control. A shipment milestone matters because it may delay a batch release, consume remaining shelf life, trigger temperature excursion review, miss a hospital delivery window, or force an inventory reallocation. A supplier delay matters because it may affect production sequencing, safety stock, alternate sourcing, and customer allocation. The control tower's job is to connect those relationships early enough for teams to act.

It should bring together product, batch, shelf life, serialization status, quality release, inventory, demand, open orders, supplier and CMO commitments, production constraints, shipment milestones, predictive ETA, carrier performance, lane reliability, temperature signals, dwell time, port and weather risk, and the cost, service, and working-capital exposure attached to each.

The key is not more data. It is a live operating picture that preserves context.

A pharma control tower should answer six questions without forcing teams into a cross-functional war room every time:

What changed?

Which internal or external signal moved materially: supplier commitment, production output, inventory position, quality release, shipment ETA, temperature status, lane condition, demand priority, or partner response?

What will be affected?

Which batches, SKUs, DCs, markets, customers, orders, patient-critical commitments, and regulatory constraints now sit inside the impact path?

How much time remains?

How many hours or days are available before the issue becomes irreversible or materially more expensive to mitigate?

How much value is exposed?

What is the expected service, revenue, inventory, expiry, quality, cost, and compliance impact if no action is taken?

Which responses are feasible?

Can the business reallocate inventory, expedite transport, switch mode, pull from another node, reprioritize production, use alternate supply, change release sequence, split orders, or communicate controlled service impact?

Who must approve the decision?

Which decision requires planner approval, quality approval, logistics approval, commercial sign-off, finance review, regulatory input, or leadership escalation?

The control tower earns its place when it moves teams from “where is the issue?” to “what decision is required, by whom, and by when?”

RISK INTELLIGENCE

Risk intelligence is the layer that decides what deserves attention.

A pharma supply chain produces thousands of signals every day. Most are noise. Some are operationally inconvenient. A few change service, compliance, patient availability, cost, or revenue outcomes. The role of risk intelligence is to separate those categories before teams waste time on the wrong problems.

- A delayed shipment only matters when the network has no room to absorb it. If the product is already low-cover, tied to a priority order, sitting inside a short shelf-life window, or has no practical alternate stock, the delay becomes a business risk.
- Low stock is not the same as lost service. The risk emerges when projected demand will burn through available inventory before replenishment, quality release, or a feasible node transfer can recover the position.
- Cold-chain alerts need context before they become decisions. The exposure depends on the excursion profile, product sensitivity, release rules, replacement lead time, and whether the affected shipment is tied to a customer or market that cannot wait.

This is why risk intelligence must work top-down: business impact first, then risk type, then signal, then event cascade. The most mature models do not simply detect exceptions. They translate exceptions into exposure.

A useful risk model should score:

Time to impact

When will the issue hit service, production, expiry, quality, or cost?

Value at risk

What revenue, order value, inventory value, margin, penalty, or working capital is exposed?

Service exposure

Which OTIF, fill-rate, allocation, hospital, distributor, or customer commitments are likely to fail?

Expiry and shelf-life exposure

Will the response create excess, aging, write-off, or market-ineligible stock?

Quality and compliance dependency

Is the mitigation blocked by batch release, deviation review, serialization, documentation, validation, or market-specific rules?

Probability of materialisation

How likely is the risk to become real, based on current signals, historical partner performance, lane reliability, and remaining response time?

Cost to mitigate

What is the cost of each response: premium freight, production resequencing, alternate supply, split shipment, transfer, write-off risk, or customer concession?

Downstream dependency

Which production lines, DCs, orders, markets, customers, and partner commitments will be affected next?

Decision owner and approval path

Who can act, who must approve, and which decisions require escalation?

The goal is not to predict every disruption. It is to prioritize the few decisions that materially change the outcome.

Risk intelligence also needs memory. If a mitigation was accepted but failed, the model should learn. If a planner rejected a recommendation because quality release would not clear in time, the model should learn. If a supplier repeatedly commits late on the same material family, the scoring logic should learn. If a lane looks reliable on average but fails during specific weather, port, or customs patterns, the model should learn.

This is how the operating model moves from reactive exception management to controlled resilience.

AGENTIC AI

Agentic AI should increase decision speed without reducing accountability.

In pharma, agents should not replace governance. They should do the assembly work that slows governance down.

The highest-value agentic use cases are the tasks that sit between signal detection and decision approval: reading unstructured partner updates, matching external events to internal flows, identifying affected batches and orders, preparing root-cause context, comparing mitigation options, drafting recommended action packages, routing approvals, and capturing feedback from accepted, modified, or rejected decisions.

This is where agents change the operating model.

An analyst can detect an exception. An agent can assemble the decision file.

For example, when a CMO signals a delayed batch, an agentic workflow can identify the affected SKU-market combinations, check available inventory and shelf life across nodes, compare open orders and demand priorities, assess quality release dependencies, calculate service and revenue exposure, test feasible reallocations, estimate expedite cost, and prepare the approval path. The decision still sits with the accountable owner. But the owner is no longer waiting for five teams to assemble the facts.

Different agents play different roles in the ecosystem:

- Signal agents detect anomalies across demand, inventory, supplier performance, logistics, production, quality, and external events.
- Context agents connect the signal to the digital twin: affected SKUs, batches, orders, nodes, lanes, suppliers, CMOs, 3PLs, and customers.
- Risk agents calculate exposure, probability, time to impact, and downstream dependency.
- Mitigation agents compare action packages: inter-DC transfer, alternate supplier, air switch, mode change, production resequencing, inventory reallocation, order prioritization, or customer communication.
- Workflow agents route the recommendation through the right approval path and push approved decisions into execution systems.
- Learning agents capture outcomes, planner feedback, partner behavior, and actual impact so the next recommendation is sharper.

The value is not “autonomous AI” in the abstract. The value is governed acceleration. Agents reduce the manual work required to reach a decision, while keeping the decision explainable, auditable, and owned.

DECISION GOVERNANCE

Agentic AI only works in pharma if the operating model is clear about where autonomy ends and accountability begins.

This is why decision governance cannot be treated as an afterthought. It defines which recommendations can move automatically, which need functional review, and which require formal escalation.

In pharma, most recovery decisions create a second consequence somewhere else.

Expediting a shipment may protect the order, but someone still has to justify the premium freight. Moving stock from one market may solve today's shortage and create next week's escalation elsewhere. Pulling a batch forward in the schedule may protect a customer commitment, but only if the change does not disrupt campaign efficiency or release timelines.

Some mitigations look strong on paper and fail inside the control model. A batch that has not cleared quality cannot be treated as available stock. A customer allocation decision cannot be made by supply chain alone when commercial priority, regulatory obligation, and patient impact are all in play.

Governance should define:

Decision thresholds

Which risks can be resolved automatically, which require functional approval, and which require leadership escalation?

Approval paths

Which decisions need planning, quality, logistics, finance, regulatory, commercial, supplier, CMO, or 3PL approval?

Response playbooks

Which mitigations are allowed by product type, market, temperature requirement, shelf-life window, service class, and compliance rule?

Auditability

What was recommended, why it was recommended, who approved it, what changed during execution, and what the final outcome was.

Feedback discipline

Whether the action was accepted, modified, rejected, delayed, or escalated, and whether it reduced exposure or created a second-order risk.

Pharma already understands this discipline in quality and batch release. The next step is to extend the same rigour to the supply chain decisions that determine whether a batch reaches the customer on time, in condition, and with the right documentation.

PROVENANCE AND TRUST

Gartner's 2026 supply chain technology trends make an important point for pharma: agentic AI cannot be separated from trust. The same report that highlights agentic AI also calls out product provenance, the ability to verify a product's origin, movement, condition, and chain of custody across the supply chain.

For pharma, provenance is not a future-facing technology theme. It is already part of the operating reality. Every batch carries questions the business must be able to answer: where it came from, whether it was handled correctly, whether it stayed within required conditions, whether it has cleared quality release, how much shelf life remains, and whether its documentation can stand up to review.

That is why provenance has to sit inside the control model, not beside it.

A control tower should not treat stock as available unless the batch can actually be used. Risk intelligence should not score an exception without understanding release status, shelf life, temperature history, and documentation risk. Agentic AI should not recommend a mitigation unless it can show the evidence behind the recommendation. Governance should not approve a response without preserving the decision record.

This is the connection that matters. Agents should not bypass pharma controls. They should make controls easier to apply at speed.

A good agentic system should show what it saw, what it inferred, which provenance checks it applied, which constraints shaped the recommendation, which options were rejected, and what evidence supports the final decision.

That makes AI useful in a regulated environment because it strengthens the control record rather than creating another black box.

THE ADOPTION REALITY CHECK

One reality check belongs here.

Only 19% of companies in McKinsey's December 2025 survey reported deploying AI at scale, roughly flat year on year, while most remained in planning or piloting.

That gap matters. There is still a wide distance between agentic AI as a technology trend and agentic AI as an operating capability.

For pharma supply chains, the answer is not to wait for enterprise-wide autonomy. It is to start with narrow, governed, high-value applications where the decision logic is clear and the cost of delay is visible: batch-delay impact assessment, cold-chain exception triage, shortage risk prioritisation, supplier commitment monitoring, shipment recovery, inventory reallocation, and approval routing.

That is how agentic AI earns trust. Not by acting everywhere, but by helping teams make the few decisions that protect service, compliance, cost, and patient availability before the window closes.

HOW THE ECOSYSTEM WORKS TOGETHER

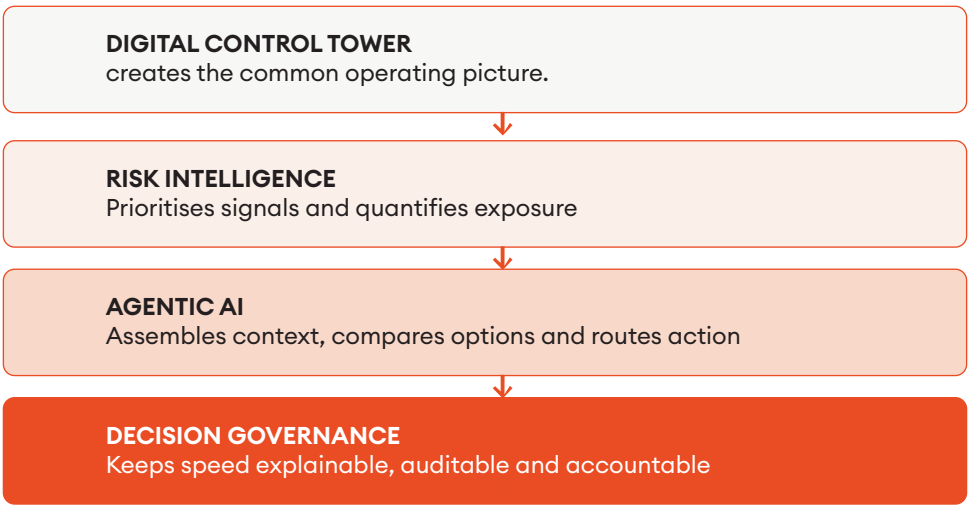


Figure: Four layers of the control model

Together, these layers convert the pharma supply chain from a monitored network into a controlled decision system.

The outcome is not just better visibility. It is fewer late decisions, fewer avoidable escalations, faster cross-functional alignment, and a more disciplined way to protect service, inventory, quality, compliance, and cost when reality moves faster than the plan.

TRANSPORT EXECUTION

Once the decision is approved, the model has to move from intelligence to execution.

This is where transport execution becomes critical.

Execution intelligence creates value only when decisions reach the physical network. Transport is where pharma risk is either contained or made expensive.

A good recommendation is not enough. The organisation must be able to tender, reroute, track, document, audit, and settle the movement. That means carrier and rate governance, route and mode optimisation, multi-leg planning, automated tendering, dock and yard coordination, predictive ETA, real-time tracking, exception escalation, proof of delivery, freight audit, settlement, and carrier performance analytics.

These capabilities matter disproportionately in pharma because transport decisions directly drive cold-chain risk, service reliability, documentation accuracy, demurrage, detention, premium freight, and customer trust.

Exhibit 3: How the control model changes the way decisions are made

Traditional mode	Control-model mode
Track shipments	Track exposed orders, products, batches, and commitments
Report aging inventory	Identify recoverable inventory before options narrow
Escalate every delay	Prioritise disruptions based on material business impact
Reroute manually	Compare feasible response options across cost, service, quality, and compliance
Review excursions after movement	Detect lane and dwell risk before product integrity is affected
Analyse freight cost after settlement	Govern premium freight, demurrage, detention, and rerouting before cost is locked in
Depend on expert firefighting	Apply governed decision playbooks with clear ownership and approval paths

The operating model is not human versus AI. It is human accountability supported by faster, better-prepared decision context.

WHERE THE VALUE SITS

Technology value in pharma supply chains is often overstated because it is presented as a software benefit. In reality, value comes from changing how decisions are made.

Two things to hold in mind when reading numbers in this space. The outcomes worth trusting are the ones where the operating model changed alongside the technology, because software alone did not produce them. And the clearest recent figures come from aerospace and consumer goods rather than pharma, so treat them as an indication of scale rather than a benchmark to plan against.

McKinsey's January 2026 analysis offers one relevant case beyond the aerospace example above. A consumer goods company that redesigned its fulfilment model around quantified sourcing risk and service-level commitments, rather than around inventory alone, reported 30% fewer stockouts while holding target service levels and freeing working capital.

Pharma's own value will sit differently, with more of it concentrated in expiry and quality-linked loss. That is the pool least visible in cross-industry benchmarks, which is a reason to measure it internally rather than wait for an external number to justify the work.

Exhibit 4: Value is created when risk is converted into a governed decision

Value pool	What is at risk	Decision that protects value
Service protection	Orders, market commitments, and customer availability	Allocate, transfer, expedite, split, or substitute
Expiry reduction	Recoverable stock approaching commercial shelf-life limits	Reallocate, reprioritise, transfer, or accelerate consumption
Cold-chain resilience	Product integrity during dwell, delay, or route disruption	Intervene, reroute, or escalate before excursion
Logistics cost control	Premium freight, demurrage, detention, and avoidable handling costs	Replan before cost is committed
Planner productivity	Time lost assembling context and triaging exceptions	Automate context preparation, prioritisation, and option comparison
Executive governance	Operational risk that has not been translated into business exposure	Prioritise action based on value, service, quality, and risk trade-offs

One mechanism runs through all six. The savings come from having fewer emergencies, not from running emergencies more cheaply. Once service has failed, a batch has expired, or freight cost has been locked in, the system can only explain the loss.

CLOSING PERSPECTIVE

The most resilient pharma supply chains will not be the ones with the most dashboards. They will be the ones that shorten the path from signal to governed action.

That path matters because pharma risk is time sensitive. Expiry risk worsens every day. Cold-chain risk worsens with dwell and delay. Port congestion becomes more expensive when teams wait for confirmation. Rerouting becomes costlier when alternatives are evaluated late. Service failures become harder to prevent once allocation windows close.

There is a timing argument as well. Digital supply chain investment slowed across industries in 2025, just as solutions capable of turning real-time risk signals into action became increasingly viable. Fewer competitors are building this capability right now than the volume of AI commentary suggests. McKinsey's own conclusion from its 2025 survey is that the companies which resume their digital agendas soonest will be best placed for the next wave of disruption.

The leadership question is no longer: Can we see the supply chain? It is: Can we control the right decision before the risk becomes irreversible?

For pharma leaders, the opportunity is not digital transformation in the abstract. It is the ability to protect product, patients, service, margin, and trust while there is still time to act.